arXiv:1511.06790v1 [astro-ph.IM] 20 Nov 2015

The LSST Data Processing Software Stack: Summer 2015 Release

Tim Jenness,[1] for the LSST Data Management Team

[1]*Large Synoptic Survey Telescope, Tucson, AZ, USA;* tjenness@lsst.org

Abstract. The Large Synoptic Survey Telescope (LSST) is an 8-m optical ground-based telescope being constructed on Cerro Pachón in Chile. LSST will survey half the sky every few nights in six optical bands. The data will be transferred to NCSA and within 60 seconds they will be reduced using difference imaging techniques and detected transients will be announced to the community in the VOEvent format. Annual data releases will be made from all the data during the 10-year mission, with unprecedented depth of coadds and time resolution of catalogs for such a large region of sky. In this paper we present the current status of the data processing software, and describe how to obtain it.

1. The Science Pipelines Software "stack"

The Large Synoptic Survey Telescope (LSST; Ivezic et al. 2008) will take about 15 TB of image data per night and after ten years of operations will have 15 petabytes of catalog data for the final data release and 0.5 exabytes of image data[1]. We are writing a suite of software packages to enable these data products to be created with sufficient quality and performance to meet the established science goals (LSST Science Collaboration 2009).

The science pipeline software will enable two key components of the data management system. The Alert Production pipelines (also known as *Level 1*) process the data from the telescope and publish alerts to the community within 60 seconds of data acquisition (Kantor 2014). Data Release Production (*Level 2*) is responsible for the annual data releases which reprocess all the data each year to generate the best possible catalogs. Both these systems will be integrated with the Calibration Products Production that continuously calculates the best calibrations for the pipelines. The software will also provide a toolkit for user-supplied code that can be used to efficiently and effectively analyze LSST data as part of *Level 3* processing or their own pipelines. Full details of the data management applications design are detailed elsewhere (Juric et al. 2016, 2013).

The LSST data management science pipelines software system, commonly referred to as the "stack", is a collection of about 40 separate packages providing functionality such as data access libraries, data models representing exposures and catalogs, source detection algorithms, astrometry fitting, and photometry and measurement algo-

[1]http://lsst.org/scientists/keynumbers

rithms. The software is written in a mixture of Python and C++[2], where the latter is used for CPU-intensive algorithms, or when the algorithms require access to complex data structures. The codebase consists of approximately 100,000 lines of Python and 110,000 lines of C++ (not including comment or blank lines and not counting expanded SWIG (see e.g. Beazley 2003) interface code).

The science pipeline packages are namespaced (with an `lsst.` root) and grouped by their functionality. The core namespaces are defined as follows:

daf The Data Access Framework is responsible for mediating between the archive resources and the application writer. The pipeline code has a completely abstract view of file I/O and only has to know how to deal with data objects representing fundamental types such as exposures and tables. Currently FITS is the internal format but the system is designed such that the internal format could be changed to HDF5, for example, and no changes would have to be made to the science pipeline code. This abstraction of the files from the code protects us against shifts in community preferences such as those discussed in Mink et al. (2015).

afw The Astronomy FrameWork provides the core classes for manipulating exposures and catalogs, including detecting sources and world coordinate handling.

ip These are the image processing classes, including packages for instrument signature removal and image differencing.

meas The measurement packages include code for determining source properties and correcting astrometry and photometry.

obs These classes provide instrument-specific knowledge to the software system, providing information to the data access framework to teach it how to interpret data from a range of optical cameras. The `obs` packages currently support data from some instruments on Subaru and CFHT, in addition to simulated LSST data. Work is ongoing to add support for DECam.

pipe Pipeline infrastructure and tasks. A task is the name for a core processing component that can be chained with other tasks to build a pipeline.

More details concerning the history behind the development of the LSST software can be found in Axelrod et al. (2010).

2. Summer 2015 release

Whilst the software is open source[3] and can be installed at any time, LSST makes formal releases of the science pipeline software at the end of each six month development cycle in the spring and autumn. The most recent release covered the summer development cycle and was labeled *Summer 2015* and released in September 2015. Detailed release notes can be found online;[4] here we provide a summary.

---

[2]Currently we support Python 2.7 but intend to also support Python 3. We are also migrating the C++ codebase to C++11.

[3]`https://github.com/lsst`

[4]`https://community.lsst.org/t/268`

Multi-band processing for coadds. New command-line tasks have been added for consistent multi-band processing of coadds. This new data processing flow carefully combines source measurements taken in multiple bands to guarantee consistent deblending across all bands, including when carrying out forced photometry, thereby enabling reliable color measurement, and ensuring that all sources are measured in each band, regardless of the bands where they are detected.

Upgraded astrometry calculation. Previously astrometry was calculated using astrometry.net code (Lang et al. 2010, ascl:1208.001) and related catalogs distributed by LSST. To improve flexibility in the code the astrometry fitter is now pluggable and now includes an alternative implementation.

Support for PSFEx. PSFEx (ascl:1301.001) is currently the state of the art external package for point spread function (PSF) determination, used in projects such as DES (Bertin 2011). LSST wrappers were created such that PSFEx could be used as a plugin in place of the built in PSF determiner.

More efficient handling of large footprints. A footprint defines the pixels associated with a particular source or blended sources. This release saw significant improvements in performance when using very large footprints.

Enable use of deblended heavy footprints in coadd forced photometry. Given the new multi-band processing for coadds we now have a reference catalog that is consistent across all bands. This catalog allows the use of the source's heavy footprints[5] to replace neighbors with noise in forced photometry, thus providing deblended forced photometry and consistent deblending across all bands. This provides much better colors for blended objects as well as measurements for drop-out objects that do not get detected in the canonical band. This functionality has been enabled for forced coadd photometry.

Significant improvements in the table class. The AFW package has a native C++ implementation of a class for manipulating table data for handling the results of detection and measurement algorithms. This release comes with some major enhancements to the internals of `afw.table` and, in particular, much better support for compound fields (such as Right Ascension/Declination tuples).

Device independent displays. DS9 (Joye 2011, ascl:0003.002) is no longer hardwired into the software and the choice of display tool is now user configurable. The intention is for the next release to include support for the Firefly visualization tool (Wu et al. 2016).

[5] A heavy footprint is a footprint that includes the pixel values.

3. Obtaining the software

The software is known to work on CentOS 6 and 7, recent Debians and Mac OS X Yosemite and Mavericks (this release is known not to work on Mac OS X El Capitan due to interactions between library path environment variables and the new System Integrity Protection feature; this has been fixed in the current development version), with a C++11 compatible compiler such as GCC 4.8.3 or later, or Apple `clang` 6 . The recommended way to install the software from source is to use the `eups` distribution installation system (Padmanabhan et al. 2015). Experimental binary releases have also been made available using a CernVM File System (CernVM-FS; Meusel et al. 2015). Full details on both these options are available on the release notes page.

Acknowledgments. The Summer 2015 release of the LSST software stack is the result of the efforts of the many people who are part of the Data Management Team at LSST, as well as outside contributors. This material is based upon work supported in part by the National Science Foundation through Cooperative Support Agreement (CSA) Award No. AST-1227061 under Governing Cooperative Agreement 1258333 managed by the Association of Universities for Research in Astronomy (AURA), and the Department of Energy under Contract No. DE-AC02-76SF00515 with the SLAC National Accelerator Laboratory. Additional LSST funding comes from private donations, grants to universities, and in-kind support from LSSTC Institutional Members.

References

Axelrod, T., Kantor, J., Lupton, R. H., & Pierfederici, F. 2010, in Software and Cyberinfrastructure for Astronomy, edited by N. M. Radziwill, & A. Bridger, vol. 7740 of Proc. SPIE, 15

Beazley, D. M. 2003, Future Generation Computer Systems, 19, 599

Bertin, E. 2011, in Astronomical Data Analysis Software and Systems XX, edited by I. N. Evans, A. Accomazzi, D. J. Mink, & A. H. Rots, vol. 442 of ASP Conf. Ser., 435

Ivezic, Z., et al. 2008, ArXiv e-prints. `arXiv:0805.2366`

Joye, W. 2011, in Astronomical Data Analysis Software and Systems XX, edited by I. N. Evans, A. Accomazzi, D. J. Mink, & A. H. Rots, vol. 442 of ASP Conf. Ser., 633

Juric, M., et al. 2013, LSST Data Management Applications Design, LDM-151, `http://ls.st/LDM-151`

— 2016, in ADASS XXV, edited by N. P. F. Lorente, & K. Shortridge (San Francisco: ASP), vol. TBD of ASP Conf. Ser., TBD

Kantor, J. 2014, in The Third Hot-wiring the Transient Universe Workshop, edited by P. R. Wozniak, M. J. Graham, A. A. Mahabal, & R. Seaman, 19

Lang, D., Hogg, D. W., Mierle, K., Blanton, M., & Roweis, S. 2010, AJ, 139, 1782. `arXiv:0910.2233`

LSST Science Collaboration 2009, ArXiv e-prints. `arXiv:0912.0201`

Meusel, R., Blomer, J., Buncic, P., Ganis, G., & Heikkila, S. 2015, Journal of Physics Conference Series, 608, 012031

Mink, J., Mann, R. G., Hanisch, R., Rots, A., Seaman, R., Jenness, T., Thomas, B., & O'Mullane, W. 2015, in Astronomical Data Analysis Software and Systems XXIV, edited by A. R. Taylor, & E. Rosolowsky, vol. 495 of ASP Conf. Ser., 11. `arXiv:1411.0996`

Padmanabhan, N., Lupton, R., & Loomis, C. 2015, EUPS — a Tool to Manage Software Dependencies, `https://github.com/RobertLuptonTheGood/eups`

Wu, X., Ciardi, D., Dubois-Felsmann, G., Goldina, T., Groom, S., Ly, L., & Roby, T. 2016, in ADASS XXV, edited by N. P. F. Lorente, & K. Shortridge (San Francisco: ASP), vol. TBD of ASP Conf. Ser., TBD